Real-time Robust Manhattan Frame Estimation: Global Optimality and Applications

Kyungdon Joo, Tae-Hyun Oh, Junsik Kim and In So Kweon

Abstract—Most man-made environments, such as urban and indoor scenes, consist of a set of parallel and orthogonal planar structures. These structures are approximated by Manhattan world assumption and be referred to Manhattan Frame (MF). Given a set of inputs such as surface normals or vanishing points, we pose an MF estimation problem as a consensus set maximization that maximizes the number of inliers over the rotation search space. Conventionally this problem can be solved by a branch-and-bound framework which mathematically guarantees global optimality. However, the computational time of the conventional branch-and-bound algorithms is rather far from real-time performance. In this paper, we propose a novel bound computation method on an efficient measurement domain for MF estimation, *i.e.*, the extended Gaussian image (EGI). By relaxing the original problem, we can compute the bounds in real-time performance, while preserving global optimality. Furthermore, we quantitatively and qualitatively demonstrate the performance of the proposed method for various synthetic and real-world data. We also show the versatility of our approach through three different applications: extension to multiple MF estimation, video stabilization and line clustering.

Index Terms—Manhattan frame, rotation estimation, branch-and-bound, scene understanding, video stabilization, line clustering, vanishing point estimation.

arXiv:1605.03730v1 [cs.CV] 12 May 2016

1 Introduction

In scene understanding, we usually look the big picture (*i.e.*, a structure of the scene) first and catch details of the scene for a deeper understanding. In case of man-made environments surrounding us, the scenes have structural forms (from the layout of a city to buildings and many indoor objects such as furniture), which can be represented by a set of parallel and orthogonal planes. In the fields of computer vision and robotics, these structures are commonly approximated by Manhattan world (MW) assumption [1]. Under the MW assumption, three orthogonal directions, typically corresponding to X, Y, and Z axes, are used to represent a scene structure, which are referred to as the Manhattan Frame (MF) in previous studies [2], [3], [4]. Recent studies have proposed a variety of MF estimation methods for scene representation [5], [6], [7], [8]. In addition to scene representation, MF estimation has been utilized as a key module for various vision applications [2], [9], [10], [11], [12], [13], [14].

In order to ensure the versatile applicability to a broad range of applications, a given MF estimation is required to have two properties: *stability* and *efficiency*. For stability, MF estimation needs to be robust against noise and outliers, and be accurate enough, since it could affect the overall performance of the subsequent applications. Also, it has to be insensitive to initialization, *i.e.*, guarantee of a global optimum. As for efficiency, the computational complexity of the MF estimation has to remain reasonable. Even when its stability is guaranteed, an MF estimation with a high order complexity is undesirable for time critical applications, *e.g.*, SLAM and navigation.

In this research, we propose a robust and real-time MF estimation approach that guarantees globally optimal solution satisfying both stability and efficiency. We pose the MF estimation problem as a consensus set maximization, and solve it through a branch-and-bound (BnB) framework [15]. Typically, bound computation is one of the main factors that disturbs the computational efficiency of the BnB framework [16]. To combat this, we suggest to relax the original problem and define a new bound that can be efficiently computed on a 2D domain (*i.e.*, extended Gaussian image, EGI [17]). This allows the bounds to be computed using a few simple arithmetic operations with linear complexity, while still preserving the global optimality and the convergence property in our BnB framework. The proposed framework is illustrated in Fig. 1. Our method is quantitatively and qualitatively validated with synthetic and real data. We also demonstrate the flexibility of our method in three applications: multiple MF estimation (*a.k.a.*, a mixture of Manhattan Frames, MMF) of a scene, video stabilization and line clustering.

This paper extends our previous work published in [18]. We dedicate our effort to further analyze the behavior of the proposed algorithm in both mathematical and empirical perspective. Specifically, we introduce several relationships essential for deriving the final results (Proposition 2), and provide detailed elucidation and proofs for each lemma and proposition that supporting theoretical guarantees of the proposed approach. We compare with a few more recent MF estimation methods with both simulation and real data quantitatively. Moreover, we extend the application section (Sec. 8.3) in several aspects. We conduct a large scale experiment with 81 million measurements for MMF estimation, and the quality of video stabilization is improved. We newly provide a robust line clustering and vanishing point estimation. These demonstrate the versatility of the proposed approach in terms of scalability and stability. The implementation of this work is available online.[1] In summary, the contributions of this work are as follows:

- We propose a branch-and-bound based, real-time MF estimation. Our approach can process around $300,000$ measurements in real-time, and 81 million measurements in near 10 seconds.
- We relax the problem and present a new and efficient bound computation with linear-time complexity, while guaranteeing a globally optimal solution. Note that our method is inherently robust by definition of consensus maximization [19], [20].
- Our method has been validated through systematic experiments and real world data. To show extensibility of our method, we

- *K. Joo, T.-H. Oh, J. Kim and I. S. Kweon are with the School of Electrical and Computer Engineering, KAIST, Daejeon, Republic of Korea. E-mail: {kdjoo369, thoh.kaist.ac.kr, mibastro}@gmail.com, iskweon@kaist.ac.kr*

1. http://sites.google.com/site/kdjoocv/cvpr2016_mf

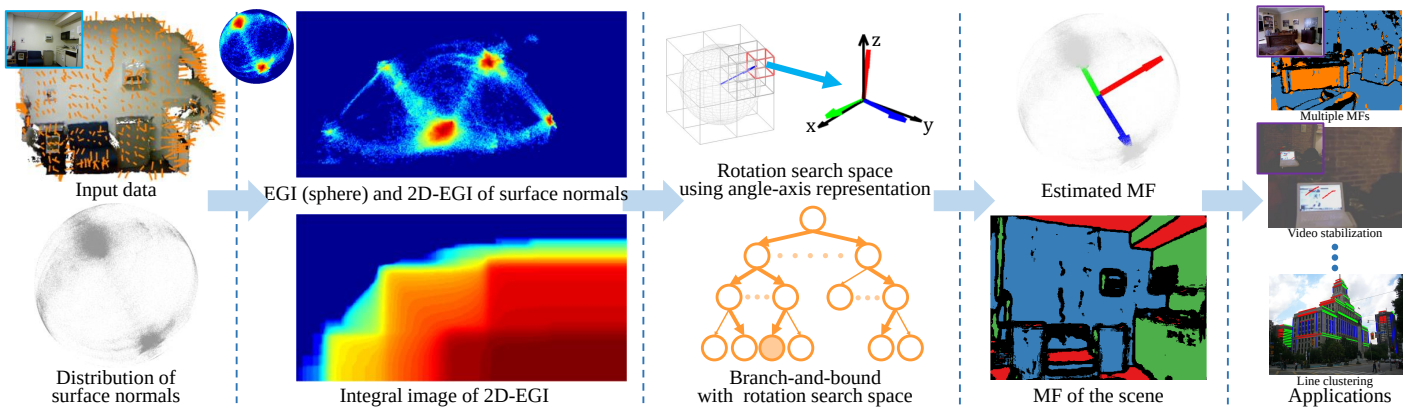

Fig. 1: Overview of the proposed MF estimation approach. *First column*: An example of input data and distribution of surface normals of a scene from the NYUv2 dataset [9]. Not limited to this, VP hypothesis from line pairs also can be fed as input without additional preprocessing. *Second column*: Its surface normal histogram, *i.e.*, EGI, the 2D-EGI and its integral image to efficiently calculate bounds. *Third column*: Illustration of the efficient bound based BnB framework using rotation search space. *Fourth column*: The estimated globally optimal MF. *Last column*: Three different applications: Multiple MFs, video stabilization and line clustering.

$\mathcal{N} = \{\mathbf{n}_i\}_{i=1}^{N}$	A set of surface normals, where $\mathbf{n}_i \in \mathbb{R}^{3\times1}$
$\mathcal{E} = \{\mathbf{e}_j\}_{j=1}^{6}$	A set of six canonical vectors
$\mathbf{R} \in SO(3)$	A 3×3 rotation matrix
$\boldsymbol{\beta} \in \mathbb{R}^{3\times1}$	Angle-axis parameterization of rotation matrix
B_π	A ball of radius π
D_{init}	An initial cube that tightly encloses the ball B_π
$\mathcal{X}_L$, $\mathcal{X}_U$	The boundary point sets of lower and upper inlier regions on the unit sphere manifold
$\hat{\mathcal{X}}_L$, $\hat{\mathcal{X}}_U$	The transferred point sets of $\mathcal{X}_L$, $\mathcal{X}_U$ on 2D-EGI
$\texttt{EGI} \in \mathbb{R}_{+}^{m\times n}$	2D-EGI representation
s	2D-EGI resolution parameter
$\angle(\cdot,\cdot)$	Angle distance between two vectors on vector space
$\phi(\cdot,\cdot)$	Angle distance along the elevation axis of 2D-EGI
$\theta(\cdot,\cdot)$	Angle distance along the azimuth axis of 2D-EGI
τ	The inlier threshold for the original problem
τ_{el}, τ_{az}	The inlier thresholds for elevation and azimuth axis
L_B, U_B	Lower/upper bound for the original problem in Eq. (2)
L_R, U_R	Lower/upper bound for the proposed relaxed problem

TABLE 1: Summary of notations.

present multiple MF estimation, video stabilization and line clustering as potential applications.

The rest of the paper is organized as follows. The next section dedicates to related work of the MF estimation. In Sec. 3, we formally define the MF estimation problem as a cardinality maximization problem. Based on this formulation, we elucidate a BnB approach most related to our goal through Sec. 4, in order to clarify its limitations and differences with our proposed method presented in Sec. 5. Then, we lead to analysis of the proposed method in Sec. 6. All the detail proofs up to Sec. 6 are placed in Sec. 10 for smooth flow of the text. In addition, the implementation details are described in Sec. 7. Then, we demonstrate the superiority and versatility of the proposed method with synthetic and real experiments and various applications in Sec. 8. We finally discuss remaining issues and conclude this paper in Sec. 9. To be concise, we summarize frequently used notations in Table 1.

2 Related Work

MF estimation of structured environments under MW assumption [1] is an essential part in high level vision tasks *e.g.*, scene understanding [2], [9], [12], layout estimation [7], [21], [22], [23], 3D reconstruction [5], [6], focal length estimation [2], [13], *etc*. Depending upon the type of input, understanding man-made scene structures (MF estimation) can be boiled down to two ways: estimating orthogonal vanishing points (VPs) in the image domain or estimating three dominant orthogonal directions (*i.e.*, rotation matrix) on 3D domain, such as depth or 3D point clouds.

Image-based approaches fundamentally rely on the perspective cues of the camera (*i.e.*, a geometric relation between VPs extracted from line segments and Manhattan structure) [5], [7], [11], [24]. There are some works that exploit additional information like structured learning [7], [23], [25] for robustness. We refer the reader to Ghanem *et al*. [3] for more detail review about image-based methods. On the other hand, recent studies have focused on accurately estimating dominant orthogonal directions of the scene on 3D using popular Kinect sensor [9], [12], [26]. As long as the related approaches are designed to reveal the Manhattan structure of a scene, we will refer to them as MF estimation approaches for the sake of clarity.

With 3D information, Silberman *et al*. [9] generate all possible MF hypotheses of a scene from both VPs of the image and surface normals. Then, all the hypotheses are exhaustively scored by the number of measurements that support a MF, and then the best one is taken. Note that even though the method exhaustively searches all hypotheses, it does not mean that a whole solution space is searched. For a room layout of indoor scene (*i.e.*, MF), Taylor *et al*. [26] estimate a gravity vector, which corresponds to the floor normal of the ground plane. This approach is based on physical assumptions that the vertical axis should be aligned with the gravity vector and large planar regions of the ground plane should be placed near the bottom of the image. Then, the other orthogonal normal vectors are sequentially estimated from the wall plane segments. Depending upon the task, occasionally, the gravity vector only required. Zhang *et al*. [27] estimate a floor normal by RANSAC [19] on a set of 3D points whose normals are close to the up-direction. Similarly to Taylor *et al*., Gupta *et al*. [12] initialize a gravity vector by the y-axis of a RGB-D image and refine it using a simple optimization process. Ghanem *et al*. [3] have recently proposed a non-convex MF estimation exploiting the inherent sparsity of the membership of each measured normal vector. Unfortunately, all the algorithms mentioned above are sub-optimal.

As the importance of the MF itself has grown, high stability and efficiency of MF estimation are desirable for general purposes. To guarantee stability, Bazin *et al*. [20], [28] propose an orthogonal VP estimation based on a BnB framework, which guarantees the globally optimal solution. Bazin *et al*. [28] present an interval analysis-based BnB framework. This strategy is improved in [20] by exploiting the rotation search space [29]. However, BnB frameworks are usually too slow for real-time applications [16]. To alleviate this limitation,

Parra *et al*. [16] propose a fast BnB rotation search method that uses an efficient bound function computation, which can register up to 1000 points in 2 seconds. However, it is still inadequate for real-time applications. To cope with it, Straub *et al*. [4] propose a GPU-supported MF inference method that operates in real-time, but does not guarantee global optimality. In contrast to the previous studies, the proposed MF estimation method with BnB framework guarantees a globally optimal solution, as well as a real-time efficiency.

Recently, Straub *et al*. [2] suggest a new perspective on MF, which they firstly call a mixture of Manhattan Frames (MMF). It is a more flexible representation based on a motivation that strict MW assumption cannot represent many real-world man-made scenes. Inspired by this work, we also extend our method to multiple MF estimation as an application.

3 Problem Statement

The type of input for our problem can either be 3D surface normals (from depth map or 3D point cloud) or VP candidates (from line segments). For simplicity, we only consider surface normal as the input type in the main context, but we present the application with VP, *i.e*., line clustering in Sec. 8.3 later. Given a set of surface normals $\mathcal{N} = \{\mathbf{n}_i\}_{i=1}^{N}$, our goal is to estimate a MF of a scene, which consists of three orthogonal directions.

As implied by the orthogonal property of MF, the process of MF estimation is equivalent to estimating a proper rotation matrix $\mathbf{R} \in SO(3)$ which transforms the standard basis of a coordinate to new three orthogonal basis best aligned with dominant surface normals up to sign. Since a direction vector and its flipping vector indicate the same structural support, we incorporate both the basis of a coordinate and its flipping vectors into a set $\mathcal{E} = \{\mathbf{e}_j\}_{j=1}^{6}$[2] of six canonical vectors. Then, estimating the optimal rotation matrix $\mathbf{R}^*$ with respect to the number of inliers can be formulated as the following cardinality maximization problem:

$$\underset{\mathbf{R}\in SO(3)}{\arg\max} \sum_{i=1}^{N} \sum_{j=1}^{6} [\![\angle(\mathbf{n}_i, \mathbf{R}\mathbf{e}_j) \leq \tau]\!], \tag{1}$$

where $\angle(\mathbf{a}, \mathbf{b})$ is the angle distance between the vectors $\mathbf{a}$ and $\mathbf{b}$, τ is the inlier threshold, and $[\![\cdot]\!]$ is the indicator function which outputs 1 if a statement is true, otherwise 0. Specifically, the problem is to find the optimal rotation on the rotation manifold (*i.e*., *solution search space*) by counting the number of inlier normals in the input set (*i.e*., *measurement space*). Unfortunately, Eq. (1) is intractable to solve directly in a numerical optimization (known to be NP-hard) [28], [30]. Following the approach of Li [30], auxiliary variables $y_{ij} \in \{0, 1\}$ are introduced for indicating whether the i-th surface normal $\mathbf{n}_i$ is an inlier (y_{ij}=1) or an outlier (y_{ij}=0) for the j-th direction vector $\mathbf{R}\mathbf{e}_j$. This leads to an equivalent integer programming problem of Eq. (1) as:

$$\begin{aligned} \underset{\{y_{ij}\},\mathbf{R}\in SO(3)}{\arg\max} \quad & \sum_{i=1}^{N} \sum_{j=1}^{6} y_{ij} \\ \text{s.t.} \quad & y_{ij} \angle(\mathbf{n}_i, \mathbf{R}\mathbf{e}_j) \leq y_{ij}\tau, \\ & y_{ij} \in \{0, 1\}, \forall i = 1 \cdots N,\ j = \{1 \cdots 6\}. \end{aligned} \tag{2}$$

Solving Eq. (2) is still challenging due to the non-linear geodesic distance measure, the non-linearity of the rotation manifold parameterization and two families of unknowns engaged in a non-linear way. Although it constitutes a challenging type of non-convex problem, it can be dealt with the BnB framework described in Sec. 4.

2. *i.e*., $\mathbf{e}_1 = [1\ 0\ 0]^\top$, $\mathbf{e}_2 = [0\ 1\ 0]^\top$, $\mathbf{e}_3 = [0\ 0\ 1]^\top$, $\mathbf{e}_4 = -\mathbf{e}_1$, $\mathbf{e}_5 = -\mathbf{e}_2$ and $\mathbf{e}_6 = -\mathbf{e}_3$.

4 Branch-and-Bound with Rotation Search

Branch-and-bound (BnB) is a general framework for global optimization [15], which conforms to the divide and conquer style procedure. The basic idea of BnB is to recursively divide a solution space into smaller sub-spaces[3] (*i.e*., *branch*) and test each sub-space whether it is possible to contain a global optimal solution with feasibility test. Typically, in the feasibility test, bound functions are defined such that the function outputs upper (or lower) bound of all possible objective values in a given sub-space. The feasibility test is conducted by checking values called *bound* of the given sub-space, whether the sub-spaces are infeasible. In such a way, the infeasible sub-spaces are excluded from search space at an early stage and the remaining sub-spaces are subdivided for further search until an optimal solution or a desired accuracy is reached. The keys of BnB framework are to define an appropriate search space and a bound function. In this section, we review the key parts of BnB for MF estimation and discuss its computational limitation.

4.1 Branching Part

The first key part of BnB is to define an appropriate search space, where it is a rotation space in our MF estimation problem. We employ the angle-axis parameterization to represent a rotation matrix $\mathbf{R}$, which is formed by a three-dimensional vector $\boldsymbol{\beta}$ in a ball B_π of radius π, whose direction $\boldsymbol{\beta}/\|\boldsymbol{\beta}\|$ and norm $\|\boldsymbol{\beta}\|$ specify the axis and angle of a rotation matrix $\mathbf{R}$ [29]. In the angle-axis parameterization, any rotation can be represented by a point in the ball B_π which is equivalent to the search space in this problem. Following convention, let D_{init} be an initial cube that tightly encloses the ball B_π. The cube representation makes the subdivision operation easy, because it is axis-aligned. We divide the rotation search space into smaller sub-spaces by octal subdivision of the cube for branching, as illustrated in the third column of Fig. 1.

4.2 Bounding Part

For the rotation search problem, the useful and efficient bound computation is suggested by Bazin *et al*. [20]. We directly re-state the result of Bazin *et al*. for referring to the connection with our problem in Eq. (2).

Proposition 1 (Bazin *et al*. [20])**.** *Given a cube D with the half side length σ and the rotation $\bar{\mathbf{R}}$ corresponding to the center of the cube D, the solutions of the following systems are the valid lower and upper bounds, L_B and U_B, of the inlier cardinality for any rotations in the cube D, respectively.*

$$\begin{aligned} L_B = \max_{\{y_{ij}\}} \quad & \sum_{i=1}^{N} \sum_{j=1}^{6} y_{ij} \\ \textit{s.t.} \quad & y_{ij} \angle(\mathbf{n}_i, \bar{\mathbf{R}}\mathbf{e}_j) \leq y_{ij}\tau, \\ & y_{ij} \in \{0, 1\}, \forall i = 1 \cdots N,\ j = \{1 \cdots 6\}. \end{aligned} \tag{3}$$

$$\begin{aligned} U_B = \max_{\{y_{ij}\}} \quad & \sum_{i=1}^{N} \sum_{j=1}^{6} y_{ij} \\ \textit{s.t.} \quad & y_{ij} \angle(\mathbf{n}_i, \bar{\mathbf{R}}\mathbf{e}_j) \leq y_{ij}(\tau + \sqrt{3}\sigma), \\ & y_{ij} \in \{0, 1\}, \forall i = 1 \cdots N,\ j = \{1 \cdots 6\}. \end{aligned} \tag{4}$$

Proof. Proof is deferred to [20]. □

In Proposition 1, the solutions of Eqs. (3) and (4), L_B and U_B, are simply obtained by exhaustively checking the inlier constraint for each normal with respect to the given rotation $\bar{\mathbf{R}}$ that corresponds to

3. This is irrelevant to the definition of *subspace* in linear algebra literature. Rather, it represents partitions of a solution space. We abuse this terminology throughout this paper without confusion.

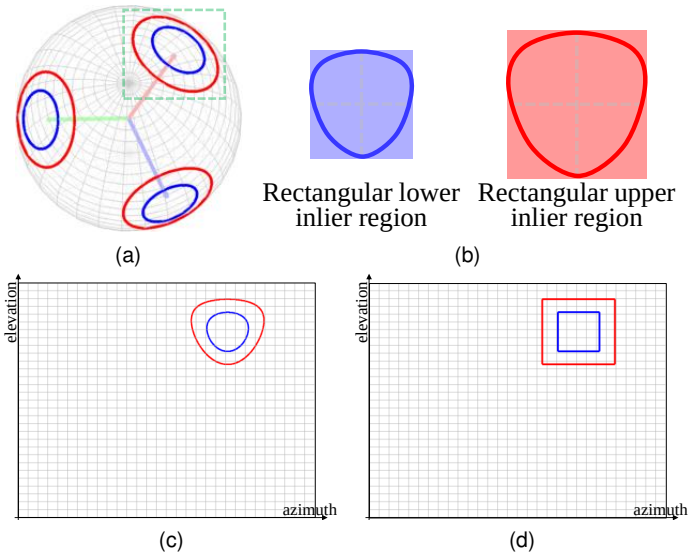

Fig. 2: Illustration of efficient inlier regions. (a) Boundary point sets $\mathcal{X}_L$ (blue) and $\mathcal{X}_U$ (red) of inlier region for lower and upper bounds of the original problem (*i.e.*, regions defined by Eqs. (3) and (4)) on a spherical domain. We visualize only three direction vectors for illustration purpose. (b) An illustration of rectangular inlier region. (c) An example of the transferred boundaries $\hat{\mathcal{X}}_L$ and $\hat{\mathcal{X}}_U$ on the 2D-EGI. (d) An example of the boundaries of the rectangular bounds defined by Eqs. (6) and (7) on the 2D-EGI.

the center of a given cube D. We refer to this method as *exhaustive BnB* to distinguish with our approach. In the above method, a single evaluation of lower and upper bounds has $O(N)$ complexity with respect to inlier computation. It is linear to the number of input normals, but along with the exponentially increased number of cubes in the branching step, this bound computation emerges as a main computational bottleneck in BnB frameworks.

5 Efficient Bound Computation

As aforementioned, the computational bottleneck arises from counting the cardinality one by one. We avoid this exhaustive counting by transferring data to an efficient measurement space which is extended Gaussian image (EGI) [17]. On EGI, we relax the original problem defined in Eq. (2) to fully exploit a structural benefit of EGI, and propose new efficient bound functions. We describe the EGI representation at the most basic level in Sec. 5.1, then we introduce the new bound functions in Sec. 5.2.

5.1 Efficent Measurement Space (2D-EGI)

For EGI representation, we discretize an unit sphere (*i.e.*, *measurement space*) by simply tessellating it along azimuth and elevation axes, which have the ranges of $[0, 360^\circ)$ and $[0, 180^\circ)$ respectively. Then, a direction vector, originated from the origin of the sphere, is approximated and represented as its elevation and azimuth angles in discrete unit. On the discretized elevation and azimuth coordinate, we construct a histogram of measurements (surface normals). As the histogram is known as an approximation to extended Gaussian image (EGI) representation [17], we will refer the EGI representation and the surface normal histogram interchangeably throughout this paper. The EGI representation on a 3D unit sphere can be directly transferred to the 2D-EGI domain (*a.k.a.* equirectangular projection), which provides powerful advantages as it allows the direct 2D representation [31] (especially with respect to memory access). The examples of EGI representation are shown as the heat map on sphere (3D-EGI) and equirectangular map (2D-EGI) in the second column of Fig. 1. The notations of EGI are formalized as follows. We denote 2D-EGI space as $\texttt{EGI} \in \mathbb{R}_+^{m \times n}$, where m and n are the height and width as $180 \times s$ and $360 \times s$ respectively, and s denotes the EGI resolution parameter meaning that each axis unit is $1/s^\circ$. There is a trade-off between accuracy and computation time by adjusting the EGI resolution.

When computing the bound values given a rotation (the center of a given cube), the inlier cardinality is counted for the normal vectors within the regions specified by a given inlier threshold (*e.g.*, τ or $\tau + \sqrt{3}\sigma$). We call this region as *inlier region*. In the case of 3D-EGI, following Proposition 1, an inlier region can be represented as a circular region as illustrated in Fig. 2a, where their boundaries of inlier regions are denoted by blue and red circles, which correspond to lower and upper bounds respectively. Let $\mathcal{X}_L$ and $\mathcal{X}_U$ be the sets of densely sampled boundary points of the lower and upper inlier regions around the six direction vectors $\{\mathbf{Re}_j\}_{j=1}^6$ on the spherical manifold of the measurement space.

Likewise, in 2D-EGI, the bound values can be computed by summing over the surface normal histogram within inlier regions, but in here the shapes of inlier regions are different from the ones on 3D. Once $\mathcal{X}_L$ and $\mathcal{X}_U$ are mapped onto the 2D-EGI, the transfered boundaries $\hat{\mathcal{X}}_L$ and $\hat{\mathcal{X}}_U$ on the 2D-EGI domain are appeared as curved shape inlier regions as illustrated in Fig. 2c. The upper and lower bounds, L_B and U_B, can be efficiently computed by summing the histogram values within the transferred boundaries on the 2D-EGI, but still it has a linear complexity with respect to the area of the curved inlier regions. We instead relax the problem, so that further speed-up is possible for the bound computation.

5.2 Rectangular Lower and Upper Bounds

Regardless of tightness of bounds, any valid bounds[4] guarantee global optimality in BnB framework with breadth-first-search strategy [15]. By slightly relaxing the inlier condition on 2D-EGI domain, we can improve computational efficiency significantly while preserving the global optimality if its bounds are still valid.

Constraint Relaxation Since the transfered boundaries $\hat{\mathcal{X}}_L$ and $\hat{\mathcal{X}}_U$ on the 2D-EGI have non-linear shapes, exhaustively traversing within the transfered boundaries is mandatory for every computation of the bounds. We instead relax the problem defined in Eq. (2) by replacing the constraint introducing non-linear shapes with new axis-aligned inlier constraints as:

$$\begin{aligned} \underset{\{y_{ij}\},\mathbf{R}\in SO(3)}{\arg\max} \quad & \sum\nolimits_{i=1}^{N}\sum\nolimits_{j=1}^{6} y_{ij} \\ \text{s.t.} \quad & y_{ij}\phi(\mathbf{n}_i, \mathbf{Re}_j) \le y_{ij}\tau_{\text{el}}, \\ & y_{ij}\theta(\mathbf{n}_i, \mathbf{Re}_j) \le y_{ij}\tau_{\text{az}}, \\ & y_{ij} \in \{0,1\}, \forall i = 1\cdots N,\ j = \{1\cdots 6\}, \end{aligned} \tag{5}$$

where $\phi(\cdot,\cdot)$ and $\theta(\cdot,\cdot)$ are the angle distances between two vectors along the elevation and the azimuth axes of EGI respectively, and τ_{el} and τ_{az} are inlier thresholds for each axis. We will discuss how to choose these inlier thresholds later. These constraints form a box constraint.

Rectangular Bounds We apply BnB to the relaxed problem in Eq. (5) by defining new valid lower and upper bound functions similar to Proposition 1. In order for new bound functions to be closest to the original problem in Eq. (2), we find the tightest circumscribed rectangles of $\hat{\mathcal{X}}_L$ and $\hat{\mathcal{X}}_U$. We call these inlier regions as the *rectangular inlier regions*, as shown in Fig. 2d. Once the boundaries is restricted to be axis-aligned and rectangle, the new (relaxed) boundaries are

4. The formal mathematical validity of bounds will be discussed later.

uniquely defined by finding the tightest circumscribed rectangle with four corner points along the elevation and azimuth axes, such that the rectangle includes the transferred point sets $\hat{\mathcal{X}}_{L,U}$ as shown in Fig. 2b. Then, the bound functions of the relaxed problem in Eq. (5) can be defined as follows:

$$\begin{aligned} L_R = \max_{\{y_{ij}\}} \quad & \sum_{i=1}^{N}\sum_{j=1}^{6} y_{ij} \\ \text{s.t.} \quad & y_{ij}\phi(\mathbf{n}_i, \bar{\mathbf{R}}\mathbf{e}_j) \le y_{ij}\tau_{\text{el}}, \\ & y_{ij}\theta(\mathbf{n}_i, \bar{\mathbf{R}}\mathbf{e}_j) \le y_{ij}\tau^{L}_{\text{az}}(\phi(\mathbf{e}_2, \bar{\mathbf{R}}\mathbf{e}_j)), \\ & y_{ij}\in\{0,1\}, \forall i=\{1\cdots N\},\ j=\{1\cdots 6\}, \end{aligned} \tag{6}$$

$$\begin{aligned} U_R = \max_{\{y_{ij}\}} \quad & \sum_{i=1}^{N}\sum_{j=1}^{6} y_{ij} \\ \text{s.t.} \quad & y_{ij}\phi(\mathbf{n}_i, \bar{\mathbf{R}}\mathbf{e}_j) \le y_{ij}(\tau_{el} + \sqrt{3}\sigma), \\ & y_{ij}\theta(\mathbf{n}_i, \bar{\mathbf{R}}\mathbf{e}_j) \le y_{ij}\tau^{U}_{\text{az}}(\phi(\mathbf{e}_2, \bar{\mathbf{R}}\mathbf{e}_j)), \\ & y_{ij}\in\{0,1\}, \forall i=\{1\cdots N\},\ j=\{1\cdots 6\}, \end{aligned} \tag{7}$$

where $\mathbf{e}_2$ is the basis of the y-axis of the 3D measurement space to measure an elevation angle, $\tau^{L}_{\text{az}}(\cdot)$ and $\tau^{U}_{\text{az}}(\cdot)$ are inlier threshold functions for the azimuth axis, which can be computed in advance by finding a circumscribed rectangle of $\hat{\mathcal{X}}_L$ and $\hat{\mathcal{X}}_U$ and be stored into a look-up table as will be specified in Sec. 7.

The feasibility constraints in the above rectangular inlier regions are interpreted as whether a pixel in the 2D-EGI map is inside the rectangle regions defined by inlier thresholds. Thus, the bound computation can be done by summing up the values in the rectangle regions, which is efficiently computed using the integral image [32], *i.e.* only with add and subtract operations of four corner values of the rectangular inlier region on integral image. The proposed rectangular bound has the following properties.

Lemma 1. *For a cube D, let c^*_D be the optimal inlier cardinality of the relaxed problem (Eq. (5)) for $\mathbf{R} \in D$, then the bounds L_R and U_R obtained by the proposed method satisfy $L_R \le c^*_D \le U_R$.*

Lemma 1 gives us the fact that the rectangular upper and lower bounds always cover the maximum cardinality of a given cube. By this fact, we can test whether it is possible for a cube to have globally optimal cardinality in a such way comparing upper bound of a cube with maximum lower bound over all the cube. However, this does not elucidate tightness of gap between the rectangular bounds, so we need the following property to see the tightness of gap at an extreme.

Lemma 2. *When the maximum half side length σ of a cube D goes to zero by sub-divisions in Alg. 1, then there exists a small $\epsilon \ge 0$ such that $U_R - L_R \le \epsilon$.*

Lemma 2 elucidates the convergence of the gap between the rectangular upper and lower bounds at an extreme. Thus, in an ideal case, we can guarantee to get point estimate as the number of iteration increases to infinity. These two lemmas imply that the proposed bound functions of Eqs. (6) and (7) are valid bounds in the perspective of BnB framework. These are indeed useful to further theoretically analyze behavior of the rectangular bound in BnB, which will be discussed in Sec. 6.

5.3 Algorithm Procedure

The BnB procedure of the proposed method is formalized in Alg. 1. The algorithm reduces volume of the search space iteratively by rejecting sub-spaces with the feasibility test until it converges to a globally optimal value or reaches a desired accuracy level. At first, the cube-list $\mathcal{L}$ is initialized with the cube D_{init} that encloses the rotation ball B_π. At every iteration, each cube in the cube-list is subdivided into octal sub-cubes with the half length size of its parent cube and stored in the cube-list, while removing the parent cubes from the list. For each sub-cube, rotation center, upper and lower bounds are computed, and then feasibility test is conducted based on the rectangular bounds (*c.f.* Sec. 5.2) as follows. From the cube-list, we discard cubes of which the upper bounds are smaller than the current maximum lower bound, as they are guaranteed not to contain globally optimal solution. This procedure carried out in every subdivision until a single cube, of which the lower bound and the upper bound are same, remains or a desired accuracy is achieved. The solution is the rotation of the cube center that has the highest cardinality, *i.e., the rotation guaranteed to be global optimum*.

Algorithm 1 BnB on the Efficient Measurement Space

Initialize the cube list $\mathcal{L}$ with D_{init} *s.t.* $B_\pi \subset D_{init}$.
repeat
  Subdivide each cube of $\mathcal{L}$ (*i.e.*, $\sigma \leftarrow \sigma/2$).
  for each cube D_i in $\mathcal{L}$ **do**
    Calculate the rotation $\mathbf{R}_{D_i}$ of the cube center.
    Compute the rectangular lower L_{R_i} and upper U_{R_i} bounds. (*c.f.* Sec. 5.2).
  end for
  $L^* = \max_i L_{R_i}$, $i^* = \arg\max_i U_{R_i}$,
  $U^* = U_{R_{i^*}}$, $\mathbf{R}^* = \mathbf{R}_{D_{i^*}}$.
  Remove all the cubes from $\mathcal{L}$ such that $U_{R_i} < L^*$.
until $\exists i$, such that $L_{R_i} = U^*$ (*i.e.*, when at least one cube whose lower bound is U^* exists) or it reaches a desired accuracy level.
Output: $\mathbf{R}^*$ (*i.e.*, the rotation matrix maximizing the number of inliers).

6 Analysis

6.1 Computational Complexity

Proposition 1 [20] tells us that each bound computation has $O(N)$ complexity. For simplicity, let C be the number of cubes that should be evaluated in the BnB framework until convergence. Then, the computation complexity of the whole procedure is $O(CN)$. This is the cases of Hartley and Kahl [29] and Bazin *et al*. [20].

In our framework, constructing 2D-EGI by accumulating surface normals and the integral image take $O(N)$ and $O(H)$, respectively, only once at the initial stage, where H denotes the number of bins of `EGI`, *i.e.*, mn. In addition, since each bound computation with the 2D-EGI representation takes $O(1)$ on the integral image, the BnB procedure on the proposed problem shows a linear complexity with respect to the number of evaluated cubes C. Thus, the overall algorithm complexity is $O(C + H + N)$ which is still linear.

We can see that the proposed method is much faster than the exhaustive BnB method [20], which has a quadratic complexity $O(CN)$. In practice, given a single depth image with resolution 640×480, N is around $300,000$ samples and C is the range of hundreds of thousands to millions (or even larger). This provides a sense of what makes the proposed algorithm real-time.

6.2 Convergence

Since it is already shown that the rectangular bounds are valid bounds by Lemmas 1 and 2, we have the following convergence result and optimal property.

Proposition 2. *The BnB algorithm with the proposed rectangular bounds and breadth-first-search is guaranteed to converge to an ϵ-optimal solution against globally optimal solution. Once it converges to a point estimate, it guarantees to converge to a global optimal solution.*

We only provide the sketch of proof here, and the detail proof can be found in Sec. 10.4. In the case of $s \to \infty$ (continuous EGI), the classic convergence result of BnB with the proposed bounds, L_R and

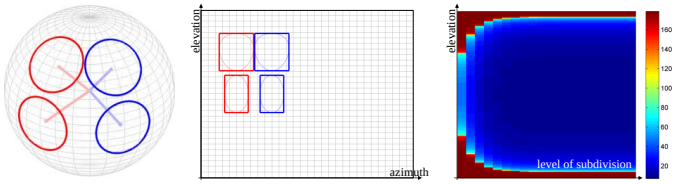

(a) (b) (c)

Fig. 3: Illustration of the properties of rectangular inlier region in the 2D-EGI domain for LUT. (a) The four direction vectors and their bound regions on sphere domain. The same color direction vectors have the same azimuth value, but different elevation values. (b) The rectangular bound regions on the 2D-EGI. Given a subdivision level, the width of rectangular region only depends on the elevation angle, while the height is fixed regardless of location. (c) The visualization of LUT for $\tau^U_{\text{az}}(\cdot)$, whose column axis indicates the elevation angle, row axis indicates the level of subdivision, and the values of entries encode the half width size of the corresponding rectangular bound region.

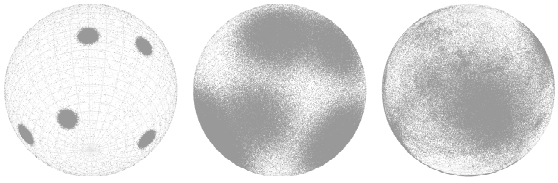

(a) $\kappa^{-1} = 0.0012$ (b) $\kappa^{-1} = 0.08$ (c) Real data

Fig. 4: Distributions of surface normals. (a) and (b) are synthetic data distribution according to κ^{-1} of vMF, 0.0012 and 0.08, respectively. (c) A sample distribution of real data from the NYUv2 dataset [9].

U_R, is valid by Lemmas 1 and 2 and the results from [15], [33]. That is, the proposed method guarantees a globally optimal solution once it converges to a point estimate. This is rather ideal case in that $s \to \infty$. Now, consider s is finite (discrete EGI). Let $\sqrt{3}(\sigma_k - \sigma_{k+1}) \leq 1/s^\circ$ be the stop criterion, where k is the level of subdivision. By the stop criterion and the relationships between the elevation (or azimuth) angle difference and the geodesic angular metric, we can prove the ϵ-optimality of the proposed method by simple arithmetics, where we can choose $\epsilon = 2/s^\circ$.

7 Implementation Details

7.1 Search Space Rejection

As aforementioned, given a rotation matrix $\mathbf{R} = [\mathbf{r}_1\ \mathbf{r}_2\ \mathbf{r}_3]$ of a cube, a direction $\mathbf{r}_i$ and its flipping direction vector $-\mathbf{r}_i$ (for $i = \{1, 2, 3\}$) indicate the same direction vector, *i.e.*, the rotation matrices $\mathbf{R}$ and $-\mathbf{R}$ indicate completely the same solution in MF estimation. Hence, since we do not need to search within the other half-sphere regions in the rotation solution space, and we exclude the half of the space.

7.2 Threshold Look-Up Table (LUT)

We observe that the shape of the rectangular boundaries on the 2D-EGI varies depending on the location of $\mathbf{r}_i$; Given a subdivision level, while the height of a rectangular boundary remains as constant,[5] its width varies according to elevation angle of $\mathbf{r}_i$ on the 2D-EGI.

Specifically, we observe a few properties regarding the shape of the rectangular inlier region (*c.f.* Sec. 5.2). Firstly, the half height of the rectangular inlier region is consistent as a constant, *i.e.*, τ and $\tau + \sqrt{3}\sigma$ for the lower and upper rectangular inlier regions, respectively. That is, τ_{el} in the 2D-EGI is equal to τ in the original domain. Secondly, the half width of rectangular inlier region only depends on the elevation angle of the mapping point of $\mathbf{r}_i$ (or $-\mathbf{r}_i$) onto the 2D-EGI. Lastly, the shapes of the rectangular inlier regions on the same elevation angle are equal, regardless of azimuth angles, as illustrated in Fig. 3b. Therefore, we generate the look-up table (LUT), which store the values of $\tau^L_{\text{az}}(\cdot)$ and $\tau^U_{\text{az}}(\cdot)$ respectively, by precomputing the half widths of the rectangular inlier region along the elevation angles and along the level of subdivision (*i.e.*, threshold $\tau + \sqrt{3}\sigma_k$) in order to reduce redundant and repetitive computation.

5. Actually, this does not hold in the polar regions of the 2D-EGI due to the range limit of the 2D-EGI map. However, by cropping the rectangle bound regions that exceed the map (or zero padding), we can equally measure the bound functions near the polar regions. This indeed provides the correct number.

For $\tau^U_{\text{az}}(\cdot)$, by virtue of the subdivision rule of rotation search space [29], we can precompute a series of bound thresholds as $\tau + \sqrt{3}\sigma_k$, where $\sigma_k = \sigma_0/2^k$ and σ_0 is the half side length of the initial cube. At the beginning of the algorithm, with the pre-computed bound thresholds, we construct an M-vector LUT for each level of subdivision, where M denotes the user specified deepest sub-division level. Entries of each vector store the calculated half width sizes of the rectangular inlier regions corresponding to the elevation angles.[6] By concatenating each vector-form LUT, we can generate a matrix form LUT (see Fig. 3c). The values of $\tau^L_{\text{az}}(\cdot)$ is stored in a vector-form LUT. By using the LUTs, we can obtain a rectangular inlier region without any computation while running the algorithm.

8 Experimental Result

In this section, we present our experimental results to explore performance of the proposed methods and to compare with other recent approaches. We systematically experiment with several criteria in simulation in Sec. 8.1, which shows how sensitive, accurate and robust the proposed algorithm is over other competitive methods. Additionally, we analyze convergence property and time profile of the proposed method. Then, in Sec. 8.2, we further evaluate our method on a real data, NYUv2 [9], which shows practicality of the proposed method in terms of accuracy and speed. Lastly, we present various extension of our core algorithm to several applications such as multiple Manhattan frame estimation, video stabilization, and line clustering and vanishing point estimation in Sec. 8.3. A large scale experiment is involved in the multiple Manhattan frame estimation application of Sec. 8.3.

8.1 Simulation

In synthetic simulations, we perform various experiments to demonstrate stability (accuracy, robustness and convergence) and efficiency (time profile) of the proposed BnB method. For synthetic data, we randomly drew three direction vectors orthogonal to each other, which correspond to ground truth MF, and additional direction vectors as outlier directions (we choose at least three outlier directions but not orthogonal to each other). We then sampled $300,000$ surface normals on von-Mises-Fisher (vMF) distribution [34],[7] which is an isotropic distribution over the unit sphere. We also uniformly sampled $20,000$ surface normals on the unit sphere to generate sparse noise (see Fig. 4). Unless specifically mentioned otherwise, we fix the inlier threshold τ as 5°, the EGI resolution parameter s as 2, and the inverse

6. The width calculation is done by finding the tightest rectangle enclosing the transfered boundaries $\hat{\mathcal{X}}_L$ or $\hat{\mathcal{X}}_U$.

7. The von Mises-Fisher distribution for the p-dimensional unit vector $\mathbf{x}$ is defined by $f_p(\mathbf{x}; \mu, \kappa) = Z_p(\kappa)\exp(\kappa\mu^T\mathbf{x})$, where μ is a mean direction, κ is a concentration and normalization constant $Z_p(\kappa) = \kappa^{p/2-1}(2\pi)^{-p/2}\mathbf{I}_{p/2-1}(\kappa)^{-1}$, where $\mathbf{I}_v$ denotes the modified Bessel function of the first kind at order v.

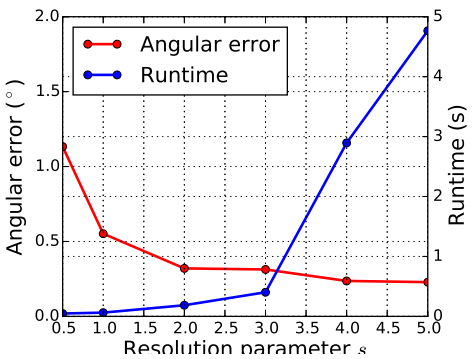

(a) Effects of EGI resolution

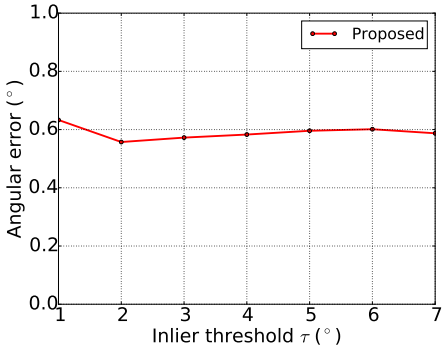

(b) Accuracy vs. inlier threshold τ

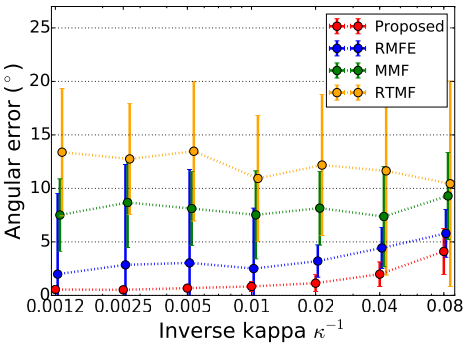

(c) Accuracy vs. noise variance

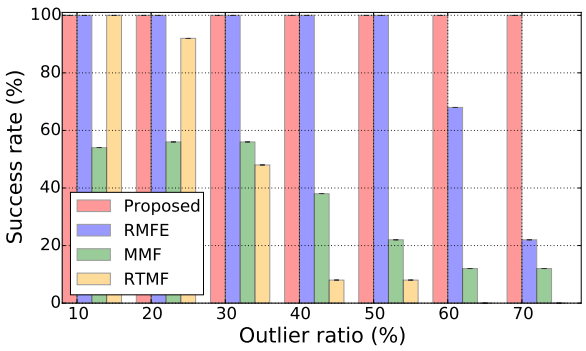

(d) Success rate against outlier ratio

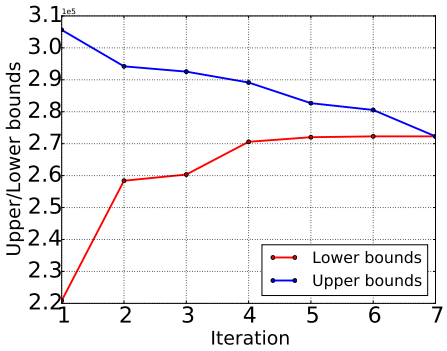

(e) Convergence graph

Fig. 5: Simulation results for observing the behaviors of the proposed method.

kappa κ^{-1} as 0.01 for each experiment. We ran each experiments 50 times on MATLAB and measured the mean of the max angular error.

Accuracy We evaluate how much accuracy of our method is affected by factors such as the EGI resolution parameter, the inlier threshold, and data distribution.

According to the EGI resolution parameter s, we test the trade-off between accuracy and runtime of the proposed method. As shown in Fig. 5a, the accuracy is improved as the resolution parameter s increases, but the ratio of the increased runtime is critical. In all other experiments, we empirically choose the resolution parameter $s = 2$ as a reasonable trade-off. Our method shows stable accuracy regardless of the inlier threshold τ (see Fig. 5b).

We also evaluate accuracy according to the variance parameter of vMF distribution, namely κ^{-1}, and compare it with recent MF estimation methods: a mixture of Manhattan Frames (MMF) [2], robust MF estimation (RMFE) [3], and real-time MF (RTMF) estimation [4].[8] For this experiment, we test κ^{-1} within a range $[0.0012, 0.08]$ on a log-scale. In Fig. 5c, each color dot indicates its mean of the maximum angular error, and each bar represents the standard deviation of maximum angular error according to κ^{-1}. MMF and RTMF estimation show relatively large errors with unstable standard deviations. RMFE shows less error than MMF and RTMF, but it also has unstable deviations due to the bias toward an outlier direction. On the other hand, the proposed method shows stable and precise accuracy, as shown in Fig. 5c.

Robustness We validate robustness of our method against outlier in a way of comparing success rates with other methods. For outlier generation, we vary a ratio of outlier surface normals from 10% to 70% for a fixed kappa value (we set κ=128), and assign outlier normals to one of aforementioned outlier directions to be clustered. We classify success of each trial if angular error against ground truth is lower than 5°. As shown in Fig. 5d, MMF and RTMF show low success rate and rapid decrease as outlier ratio increases. RMFE, on the other hand, shows stable success rate, but it is also sensitive to high outlier ratio. On the other hand, the propose method shows robustness for high outlier ratio.

8. For a fair comparison, we used the publicly available code and original parameters of each papers.

Method	MPE	MMF	ES	RMFE	RTMF	Exhaustive BnB	Proposed
θ_x	26.3°	5.3°	2.3°	2.3°	4.1°	2.9°	3.0°
θ_y	18.1°	4.6°	5.6°	4.7°	2.7°	1.8°	2.0°
θ_z	18.2°	5.3°	2.9°	2.8°	3.9°	2.8°	2.9°
Avg.	20.87°	5.07°	3.60°	3.27°	3.57°	**2.5°**	**2.63°**
Runtime (s)	2.8	148.7	21.4	0.9	-	117.06	**0.07**

TABLE 2: Comparisons of average angular error and runtime for the ground truth of the NYUv2 dataset [3].

Convergence Fig. 5e shows how many sub-divisions (iteration) are typically required until convergence. It commonly converges within 7 iterations and shows that the proposed efficient bounds are valid. Moreover, the lower and upper bounds of the proposed BnB method converge to a specific value, demonstrating the convergence property in Proposition 2 of the proposed algorithm.

Time Profiling To show improvement in the computational efficiency, we compared the time profiles of the exhaustive BnB [20] and the proposed BnB (see Fig. 6). Both methods have three steps in common: branch (subdivision), rotation center estimation and bound computation, while the proposed BnB has an additional EGI generation step for efficient bound estimation. In case of the exhaustive BnB, bound computation takes $108.178s$, which is 99.98% of the entire computational time. On the other hand, the proposed BnB takes only $4.6ms$ to compute bounds. This reduces the bound computation time by more than $20{,}000$ times, compared to that of the exhaustive BnB.

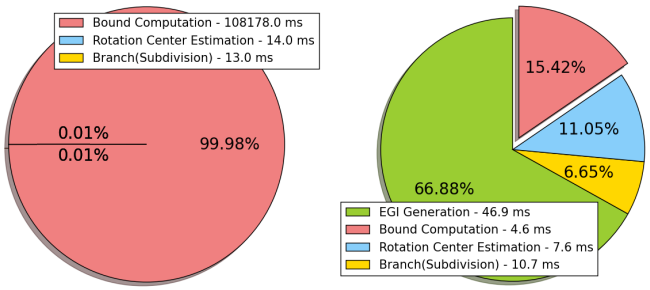

(a) Exhaustive BnB [20] (b) Proposed BnB

Fig. 6: Time profiles of the exhaustive BnB [20] and the proposed BnB approches.

8.2 Real-World Experiment

In this section, we evaluate the proposed method in terms of accuracy and speed with real-world benchmark dataset, NYUv2 [9], which contains 1449 RGB-D images of various indoor scenes. In particular, we used the recently introduced ground truth benchmark [3] on NYUv2 dataset for quantitative evaluation. We fix all the parameters of the proposed method as before. We compare the exhaustive BnB and the proposed BnB with the recent MF estimation approaches: MPE [26], MMF [2], ES [9], RMFE [3] and RTMFE [4]. In case of MMF, we run 10 iterations, calculate the angular errors for multiple cases and choose a minimum error in order to alleviate the randomness of its initialization. For RTMF, we do not compute runtime since it utilizes GPU parallelization. We test average angular errors on similar hardware configurations (*i.e.*, a 3GHz workstation on MATLAB) as done by Ghanem *et al.* [3]. For MPE, ES, and RMFE, we directly quote the results from Ghanem *et al.*

Since MPE is based on the assumption that a large portion of the scene consists of the floor plane, it is sensitive to scene clutter and outliers. MMF shows less accurate result than those of ES and RMFE. We deduce the reason for MMF's poor performance in angular errors as the absence of noise/outlier handling. ES and RMFE show comparable results, but their runtime are inefficient for real-time applications. While RTMF shows reasonable result with real-time efficieny, it heavily relies on GPU requirement. The exhaustive BnB shows the most accurate result, but its runtime is intractable in the efficiency aspect. On the other hand, the proposed BnB performs stably while achieving real-time efficiency, as shown in Table 2. The accuracy difference between the exhaustive BnB and the proposed one may come from the relaxation gap, in that we solve a relaxed problem defined in Eq. (5) which is discussed later.

8.3 Applications

From now on, we present our method to several applications: multiple Manhattan frame estimation, video stabilization, and line clustering and vanishing point estimation. These are built on the robustness of the proposed method as evaluated in the previous sections, and demonstrate practicality.

Extension to Multiple Manhattan Frames Since the strict MW assumption cannot represent general real-world indoor and urban scenes, Straub *et al.* [2] introduce a more flexible description of a scene, consisting of multiple MF, namely a mixture of Manhattan Frames (MMF). As an application, we extended the proposed method to MMF estimation by sequentially finding different MFs.

To estimate MMF, we apply a greedy-like algorithm. For a given input data, we estimate an optimal MF, and update the normal data by excluding the set of normals that corresponds to the inliers of the optimal MF. We then sequentially estimate the next optimal MF for the updated normal data. Following Straub *et al.* [2], we only regard MFs of which inlier normals are more than 15% of all valid normals as good MF estimates, to deal with poor depth measurements.

We qualitatively demonstrate our extension, *i.e.*, MMF inference for the small (indoor) and large (outdoor) scale dataset. We use NYUv2 dataset [9] as a small scale dataset. From each NYUv2 dataset, we extract around $300,000$ normal data as an input for MMF inference. The proposed MMF inference shows reasonable results, as shown in Figs. 7a and 7b.

As a large scale dataset, we use Robotic 3D scan repository[9] which freely provides 3D point clouds captured by laser scanner and pose data in various scenes. Among them, we use the point cloud captured in the city center of Bremen, Germany. This dataset consists of around $81,000,000$ 3D points. For this dataset, our method take $10.124s$ to estimate MMF in total. Out of the time, the EGI generation takes $9.765s$, which occupies most of the computational time. As shown in Fig. 7c, the proposed method properly results in 4 MFs in the scene.

Video Stabilization The goal of video stabilization is to generate a visually stable and pleasant video from a jitter video by camera shake. Depending on the information used for stabilization, the approaches can be categorized into 2D or 3D motion-based stabilization. 3D-based stabilization reflects more realistic motion information than 2D-based stabilization. Recently, Jia *et al.* [36] proposed a 3D video stabilization that exploits the rotation of camera poses obtained from a built-in gyroscope in smartphones and tablet PCs. Instead of using the 3D rotations from the gyroscope, we apply the rotation matrices obtained by our method and Kinect RGB-D data, and verify the applicability of our algorithm on video stabilization. For the experiment, we use videos of NYUv2 dataset [9]. We added rotation noise to mimic egocentric head motions, and applied to the image and depth sequences.

For visualization purpose, we overlay sampled feature trajectories, as shown in Fig. 8, which reflect the stabilization performance. The initial features are detected by FAST [38], then be tracked by KLT [39] along consecutive frames. Note that these trajectories are not used in our algorithm, but only for visualization purpose. We compare the feature trajectories of the processed videos by YouTube video editor [35], the stabilization [36] based on the estimated MF by our method, and YouTube video editor applied after stabilization with our method. Fig. 8a shows the smooth feature trajectories of the original video (ground truth) which represents the true camera motion. Fig. 8b shows the jittering feature trajectories by synthetic rotation noise applied to the original videos. Figs. 8c, 8d and 8e show the three different stabilization results for comparison.

Since YouTube video editor only uses 2D motion information, it generates smooth trajectory, but it is far from the true camera motion, while the trajectory of stabilization with our method shows similar tendency to the original feature trajectory (see Figs. 8c and 8d). This comparison shows that the 3D motion based stabilization tends to preserve the true camera motion, on the other hand, the 2D motion based stabilization only focuses on smoothing the trajectory, distorting the original camera motion. For more pleasurable result, we apply YouTube video editor to the results obtained by our 3D stabilization, so as to generate natural and smooth trajectory while preserving the true camera motion. The small difference of trajectories between Figs. 8a and 8e comes from the camera motion optimization during the 3D motion based stabilization and the warping effect of the 2D motion based stabilization. In most of the cases, our approach

9. http://kos.informatik.uni-osnabrueck.de/3Dscans/

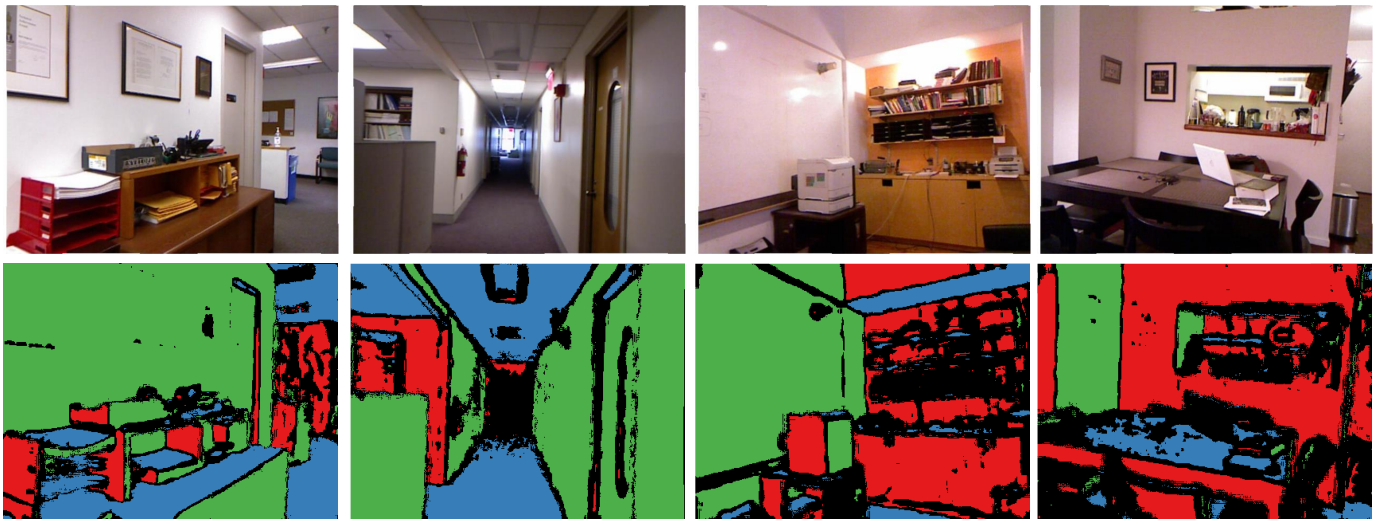

(a) Small scale - 1 MF

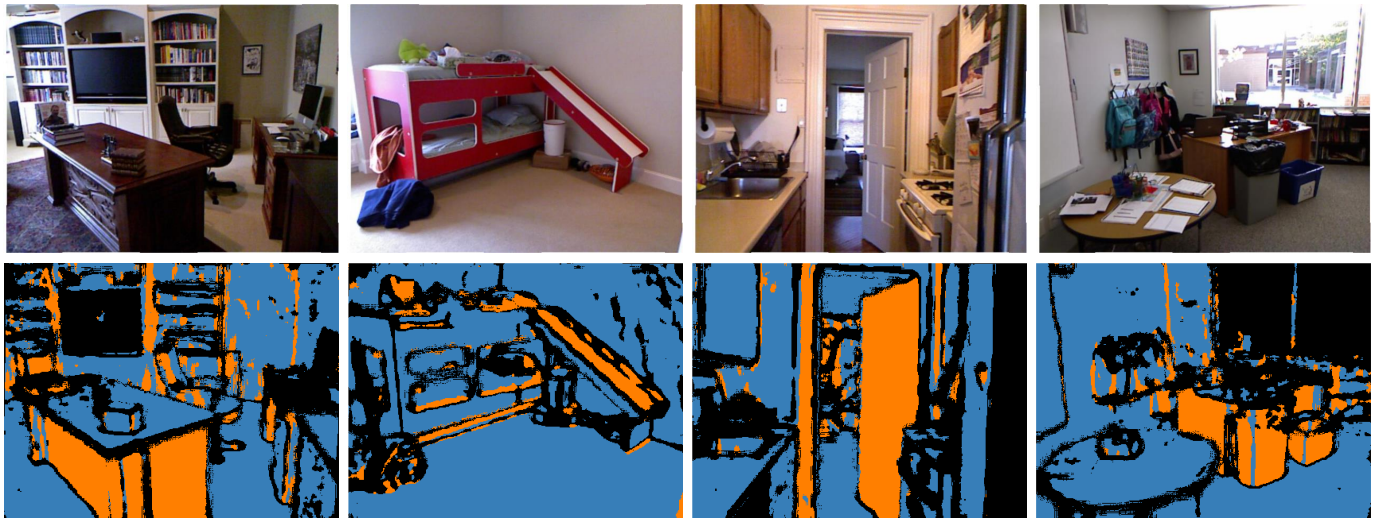

(b) Small scale - 2 MFs

(c) Large scale - 4 MFs

Fig. 7: Extension to Multiple Manhattan Frames for small and large scale dataset. [Small scale] We show the RGB images of various NYUv2 indoor scenes [9] in the odd rows, and the inferred MFs by our method in the even rows. (a) 1 MF. each color (red, green and blue) indicates MF axes. (b) 2 MFs. Each color (orange and blue) indicates each MF of the scene. [Large scale] (c) The 3D point cloud and Google earth image captured in the city center of Bremen, Germany. We estimated 4 MFs and colored each MF as red, green, blue and yellow on the Google map image. For the visualization purpose, we only crop a partial part of our results to show the inferred MMFs clearly.

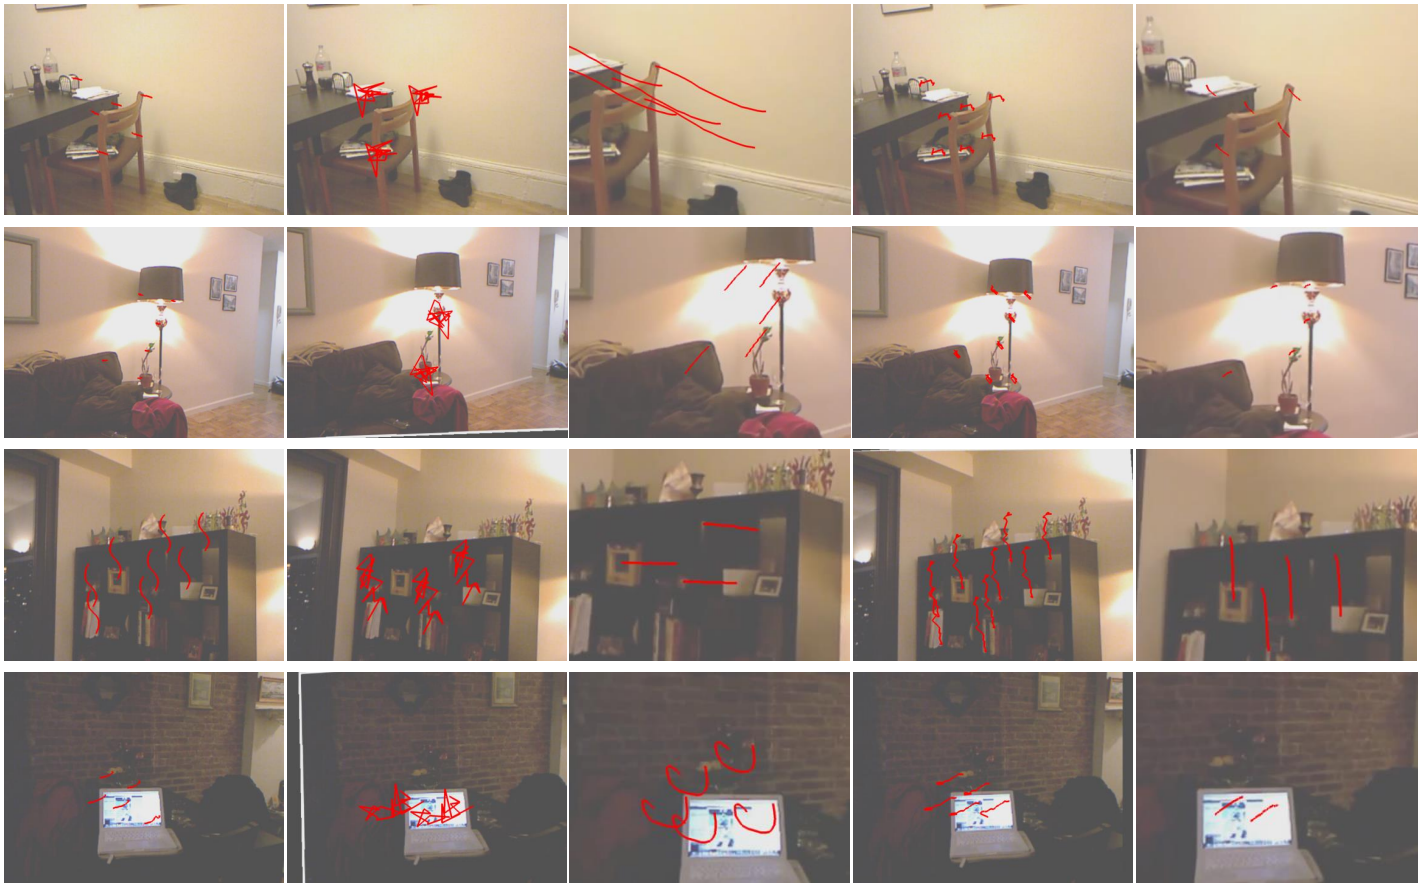

(a) Ground truth motion (b) Noise (c) YouTube video editor [35] (d) Ours (e) YouTube video editor with ours

Fig. 8: Video stabilization. Each column shows a trajectory of (a) ground truth motion, (b) motion perturbed by synthetic noise, (c) YouTube video editor [35], (d) the stabilization [36] based on our MF estimate, and (e) ours followed by YouTube video editor. Video scenarios with various camera motions are shown. Note that overlaid trajectories are only for qualitative visualization purpose. The first and last row show the stabilization comparison when the ground truth motion has noticeable translation, the second row shows the comparison when the ground truth motion is static and lastly for the third row, it shows when the ground truth motion is wavy. From the top to bottom row, the scenes are from living room 0002, living room 0005, living room 0006 and living room 0009 of the NYUv2 dataset [9].

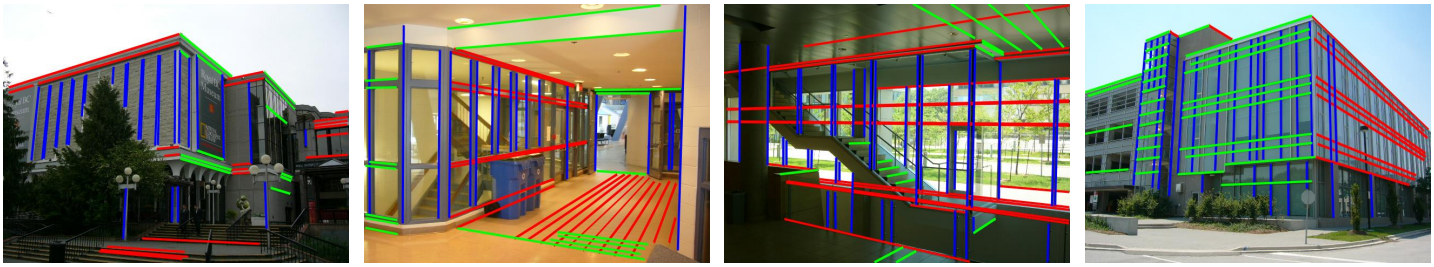

Fig. 9: Examples of line clustering obtained by our method on York urban database [37]. The same colored lines indicate a cluster associated with the same VP.

recovers the true camera motion and guides the YouTube video editor very well to generate trajectories that are very close to the true camera motion. However, there is some exceptional cases such as the third row of Fig. 8. As the true motion is a wavy motion, the true motion is flattened during the 3D camera motion optimization. Note that the depth normal obtained from the NYUv2 dataset is very noisy giving inaccurate normal information. Nonetheless, the stabilization using our method shows plausibly stabilized results.

Joint Line Clustering and Vanishing Point Estimation Given set of lines, VPs can be obtained by line clustering, and line clustering can be done by finding assignments to VPs, *i.e.*, which line belongs to which VP. Namely, it is a chicken and egg problem. In an ideal case without any image distortion and measurement error, maximum number of lines should intersect at one of VPs [28]. As input, we find a set of intersection points of all the pair of lines, and then back-project the intersection points (*i.e.*, a hypothesis direction of VPs) on the 3D Gaussian sphere using intrinsic camera parameter. This input set has the same form with surface normals on measurement space, thus VPs are represented as 3-dimensional direction vectors (*i.e.*, vanishing direction) on the sphere domain, which are orthogonal to each other and be represented as a rotation matrix $\mathbf{R}$. Therefore, the joint estimation of line clustering and VPs estimation problem can be boiled down to an MF estimation same as Eq. (1).

By applying our approach to the line clustering problem, we obtain VPs that maximize the number of intersection points that are within angular threshold from one of VPs. In camera geometry, a normal of plane induced by back-projected image line must be orthogonal to associated VP [40]. Using this relation, we simply cluster each line by assigning them to a VP they belongs to as:

$$j^* = \arg\min_j |\angle(\mathbf{n}_i, \mathbf{v}_j) - \pi/2|, \tag{8}$$

where $\mathbf{n}_i$ denotes a plane normal of i-th line segment, $\mathbf{v}_j$ denotes j-th VP for $j = \{1, 2, 3\}$, and j^* indicates the associated VP.

We use York urban database [37]. The data base contains 102 images acquired in indoor and outdoor man-made environments with intrinsic camera parameter. Each image contains a set of line segments labeled manually, vanishing points and line association to vanishing points. We apply our method to York urban dataset and compare it with ground truth. The results in Fig. 9 show that our approach successfully finds VPs and line clustering in the real world dataset.

9 Discussion and Conclusion

Based on the motivation that MF estimation requires two properties (*i.e.*, stability and efficiency) for general performance, we have presented a robust and real-time MF estimation that guarantees a globally optimal solution. This can be achieved by relaxing the original problem and computing the bounds efficiently on 2D-EGI domain. By virtue of this, the computational complexity of BnB is dramatically reduced to linear complexity. In theoretical aspects, we show the validity and convergence of the proposed bounds. Then, based on these, the convergence property to global optimality is established. To validate the stability and efficiency of our method, we demonstrate the state-of-the-art performance on various synthetic and real-world experiments by showing outperforming speed with precise accuracy and robustness. Since MF estimation is a fundamental problem, thus our proposed method, which is efficient, robust and accurate, may improve stability of other subsequent applications that adopt ours as preprocessing, including what we have shown: multiple MF estimation, video stabilization and line clustering. In the following, we discuss some open questions related to our paper:

Relaxation Gap Since we relax the original problem in Eq. (2) to Eq. (5), there is a relaxation gap between two problem. This gap leads to the accuracy difference between Exhaustive BnB and our method in Table 2. Fortunately, the difference may be negligible, but still further mathematical analysis of the relaxation gap would be a worthwhile and interesting research direction.

Further Speed-up As shown in Fig. 6, the EGI generation (mainly consist of histogram construction) is the most bottleneck for further speed-up. Obviously, the histogram computation can be easily parallelized. Interestingly, we tried GPU parallelization for EGI with surface normal data obtained from Kinect, but it does not improve the speed. This happened due to the memory transfer time from CPU memory to GPU memory. The transfer time was longer than the benefit from GPU parallelization, thus we would recommend CPU parallelization with few hundreds of thousands data. However, when the data size is large like in Fig. 7c, the parallelization will be beneficial. Note that we do not apply any parallelization for all the reported experiments in this paper.

10 Analyses of Validity, Convergence and Global Optimality of the Proposed Bounds

10.1 Preliminary for Proof

Prior to present the theoretical guarantees of the proposed algorithm, we first introduce several relationships useful for deriving the final results.

Lemma 3. *For any vector* $\mathbf{v}_1$ *and* $\mathbf{v}_2$, $\theta(\mathbf{v}_1, \mathbf{v}_2) \leq \angle(\mathbf{v}_1, \mathbf{v}_2)$.

Proof. Let $\mathbf{e}_2$ be the elevation axis (*e.g.*, y-axis in 3D space), and $\theta(\mathbf{v}_1, \mathbf{v}_2)$ is the difference angle metric along the elevation between two vectors as $\theta(\mathbf{v}_1, \mathbf{v}_2) = |\angle(\mathbf{v}_1, \mathbf{e}_2) - \angle(\mathbf{e}_2, \mathbf{v}_2)|$. Then,

$$\begin{aligned}\theta(\mathbf{v}_1, \mathbf{v}_2) =& |\angle(\mathbf{v}_1, \mathbf{e}_2) - \angle(\mathbf{e}_2, \mathbf{v}_2)| \\ \leq& |\angle(\mathbf{v}_1, \mathbf{v}_2) + \angle(\mathbf{e}_2, \mathbf{v}_2) - \angle(\mathbf{e}_2, \mathbf{v}_2)| \\ & \text{(By triangle inequality,} \\ & \angle(\mathbf{v}_1, \mathbf{e}_2) \leq \angle(\mathbf{v}_1, \mathbf{v}_2) + \angle(\mathbf{e}_2, \mathbf{v}_2)) \\ =& \angle(\mathbf{v}_1, \mathbf{v}_2).\end{aligned} \tag{9}$$

The relationship between the difference along the azimuth $\phi(\cdot, \cdot)$ and $\angle(\cdot, \cdot)$ can be defined similarly. We omit these cases. □

Lemma 4 (Lemma 1. in Hartley and Kahl [29])**.** *For any vector* $\mathbf{v}$ *and two rotations* $\mathbf{R}$ *and* $\mathbf{R}'$, $\angle(\mathbf{R}\mathbf{v}, \mathbf{R}'\mathbf{v}) \leq d_\angle(\mathbf{R}, \mathbf{R}')$, *where* $d_\angle(\mathbf{R}, \mathbf{R}')$ *is the angle lying in the range* $[0, \pi]$ *of the rotation* $\mathbf{R}'\mathbf{R}^{-1}$.

Proof. Refer to Hartley and Kahl [29]. □

Corollary 1 (Hartley and Kahl [29])**.** *Given a cube* D *and the rotation* $\bar{\mathbf{R}}$ *corresponding to the center of the cube, for any rotation* $\mathbf{R} \in D$, *we represent the radius of* r_D *of the cube in terms of the angle metric as* $r_D = \max_{\mathbf{R} \in D}(d_\angle(\bar{\mathbf{R}}, \mathbf{R}))$. *Then, we have* $r_D \leq \sqrt{3}\sigma$, *where* σ *is the half side length of the cube* D.

Proof. Refer to Hartley and Kahl [29]. □

Lemma 5. *There is one-to-one mapping between a boundary* $\mathcal{X}_{(\cdot)}$ *on the spherical measurement domain and rectangular boundary on 2D-EGI domain, up to quantization of 2D-EGI. When* $\mathcal{X}_{(\cdot)}$ *includes another* $\mathcal{X}'_{(\cdot)}$ *on spherical domain, then so it is for the respective corresponding rectangular boundaries.*

Proof. We present the basic high-level proof, because it is simple to prove the relationship by its definitions. We derive a relationship between the boundary $\mathcal{X}_{(\cdot)}$ on the spherical measurement domain[10] and the rectangular boundary.

Once the inlier threshold constraints in Eqs. (6) and (7) are specified, and equally corresponding inlier boundaries (rectangular boundaries in this case) are specified, the bound computation is nothing but finding whether y_i is 1 or 0 by checking whether the i-th data is located inside of the inlier boundaries or not. For simplicity, we only consider a boundary corresponding a rotation axis $\mathbf{R}\mathbf{e}_j$, but considering the other boundaries of the rotation axes, *i.e.*, $\{\mathbf{R}\mathbf{e}_j\}$ for $j = \{1, \cdots, 6\}$, is also analogous. A rectangular boundary on 2D-EGI is defined as a rectangle by mapping a boundary $\mathcal{X}_{(\cdot)}$ from the spherical measurement domain onto 2D-EGI, and it is uniquely defined by finding the tightest rectangle enclosing the mapped nonlinear shape boundary $\hat{\mathcal{X}}_{(\cdot)}$. Thus, a rectangular boundary has a corresponding boundary $\mathcal{X}_{(\cdot)}$.

10. These boundaries are defined by Eqs. (3) and (4)

Directly showing the opposite direction is rather difficult. Instead, we use homeomorphism property of the mapping to elucidate the opposite mapping from a rectangular boundary to a boundary $\mathcal{X}_{(\cdot)}$. Essentially, the map between spherical domain and 2D-EGI is identical to the equirectangular projection, which is a bijection mapping and preserves geometric topology (*i.e.*, homeomorphism). It means that there is a one-to-one mapping between a boundary $\mathcal{X}_{(\cdot)}$ and its corresponding rectangular boundary up to quantization of 2D-EGI (refer to Fig. 2) as well as geometrical relationships (*e.g.*, inclusion) are preserved. □

Remark 1. *Lemma 5 shows the relationship between the boundary $\mathcal{X}_{(\cdot)}$ by Eqs. (3) and (4) and the proposed rectangular boundary, which allows us to analyze the property of the rectangular boundary by $\mathcal{X}_{(\cdot)}$. Thus, we can prove related statements to rectangular bounds indirectly by $\mathcal{X}_{(\cdot)}$ yet exactly instead of using the rectangular boundary.*

10.2 Proof of Lemma 1

Proof. Let $\mathbf{R}^*$ be the optimal rotation for a cube D providing c_D^*. Since c_D^* is the optimum for all $\mathbf{R} \in D$, there is no rotation that provides larger than c_D^* in D and inlier cardinality of any rotation $\mathbf{R} \in D$ is smaller or equal to c_D^*, *i.e.*, $L_R \leq c_D^*$.

The validity of upper bound U_R can be proved using upper bound inclusion (explained soon) and Lemma 5 (homeomorphism). The validity of U_B *w.r.t.* Eq. (2) has been proved by Bazin *et al.* [20]. Given $\mathbf{R} \in D$, considering any boundary defined by Eq. (2) in D, such boundary is always included in the region of $\mathcal{X}_U$ for a center rotation $\bar{\mathbf{R}}$ of D (called as *upper boundary inclusion*). This is because any inlier of $\mathbf{R}$ *w.r.t.* the inlier threshold τ is also an inlier of $\bar{\mathbf{R}}$ *w.r.t.* $\tau + \sqrt{3}\sigma$.

Analogously, if the upper boundary inclusion is satisfied for the rectangular boundary, we have the same inlier match property, *i.e.*, upper bound U_R is valid. Since the rectangular boundary is obtained by equi-rectangular projection which satisfies homeomorphism, the upper boundary inclusion also holds for the rectangular bounds and now we have *i.e.*, $c_D^* \leq U_R$. □

10.3 Proof of Lemma 2

Proof. we now show the convergence of the bound stated as

$$\forall \epsilon \text{ and } \forall D \subseteq D_0, \exists \delta \quad \text{s.t. size}(D) \leq \delta \Rightarrow U_R^{D,k} - L_R^{D,k} \leq \epsilon, \tag{10}$$

where D_0 denotes the initial cube enclosing the solution search space, $U_R^{D,k}$ denotes the upper bound obtained by the system in Eq. (7) at iteration k with the half side length σ_k of the cube D (analogously $L_R^{D,k}$ denotes the lower bound obtained by Eq. (6)).

To easily find the ϵ bound of the most right term in Eq. (10), we use Lemma 5 again. With this fact and by definition, when the upper boundary $\mathcal{X}_U$ shrinks and gets close to the lower boundary $\mathcal{X}_L$ as reducing σ_k by subdivision, its upper rectangular boundary corresponding to $\mathcal{X}_U$ also shrinks and gets close to the lower rectangular boundary corresponding to $\mathcal{X}_L$. Thus, this relationship allows us to prove Eq. (10) by instead showing the existence of ϵ' such that $U_B^{D,k} - L_B^{D,k} \leq \epsilon'$.

As aforementioned, the inlier boundaries on the spherical domain are determined from the systems in Eqs. (3) and (4) as follows.

Given a rotation $\bar{\mathbf{R}}$ corresponding to the center of the cube D,

$$\angle(\mathbf{n}_i, \bar{\mathbf{R}}\mathbf{e}_j) \leq \tau \text{ for the lower bound.} \tag{11}$$

$$\angle(\mathbf{n}_i, \bar{\mathbf{R}}\mathbf{e}_j) \leq \tau + \sqrt{3}\sigma_k \text{ for the upper bound.} \tag{12}$$

Since the upper bound is larger than or equal to inlier cardinality for any other rotation in D, it is also true for a child cube $D' \subseteq D$. This induces $U_B^{D',k} \leq U_B^{D,k-1}$ (called as shrinkage of the upper bound). Therefore, the gap is bounded and reduced as $U_B^{D',k} - L_B^{D',k} \leq U_B^{D,k-1} - L_B^{D',k}$.

Along with this bounded gap, we have that the area difference between the lower and upper bound regions from Eqs. (11) and (12) is proportional to $(\tau+\sqrt{3}\sigma_k)-\tau = \sqrt{3}\sigma_k = \frac{\sqrt{3}\sigma_0}{2^k}$ (by the subdivision rule of the algorithm, we have $\sigma_k = \frac{\sigma_0}{2^k}$). It means that the area difference between the bound regions shrinks according to increasing k. For large k, also the divided cubes at the k-iteration must be small size by Corollary 1 (*i.e.*, $\text{size}(D) \leq \frac{\sqrt{3}\sigma_0}{2^k}$).

Consequently, as increasing k, $\text{size}(D)$ is converged to zero and the gap $(U_B - L_B)$ is bounded and reduced, then the gap $(U_R - L_R)$ of rectangular bound is also bounded and reduced (*i.e.*, $U_R - L_R \leq \epsilon$) by Lemma 5, which concludes the proof. □

10.4 Proof of Proposition 2.

Proof. Firstly, suppose a case of $s \to \infty$, *i.e.*, the measurement space (EGI) is continuous, on which the classic convergence result of BnB is given. Since the BnB with breadth-first-search is guaranteed to converge to the globally optimal solution, when the lower and upper bounds are valid [15], [33], once all the bounds L_R and U_R are valid, the proposed BnB framework satisfies the statement for global optimality. We already see that L_R and U_R are valid bounds by Lemma 1. By the results in [15], [33], the BnB algorithm with the proposed bounds in Eqs. (6) and (7) is guaranteed to converge to a globally optimal solution once it converges to a point estimate.

Consider s is finite, *i.e.*, EGI is a discretized domain as described in this work. On the discrete domain, there is a certain sub-division level that no longer reduces the gap between the subsequent upper bound values, *i.e.*, $|U_R^{k+1} - U_R^k|$ due to the resolution of EGI. We can choose the stop criterion ξ by this fact as the difference of the thresholds between the subsequent upper bound systems.

$$\xi_k = (\tau_{\text{el}} + \sqrt{3}\sigma_k) - (\tau_{\text{el}} + \sqrt{3}\sigma_{k+1}) = \sqrt{3}(\sigma_k - \sigma_{k+1}), \tag{13}$$

where k is the number of iteration (equal to the level of subdivision). For simplicity, we prove the convergence behavior only for the elevation axis in this proof, because we can separately consider the elevation and azimuth axis thresholds (we can stop the algorithm when both or either stop criteria is satisfied). However, the case of the azimuth axis can be proved in the same way.

Since we have $\sigma_k = \frac{\sigma_0}{2^k}$ by the sub-division rule, where σ_0 denotes the half-side length of the initial cube, we can rewrite $\xi_k = \frac{\sqrt{3}\sigma_k}{2}$. Due to the resolution limit by $1/s^\circ$, when $\xi_k \leq 1/s^\circ$, more branching does not improve the accuracy. Thus, we have the following criterion:

$$\xi_k = \frac{\sqrt{3}\sigma_k}{2} \leq \frac{1}{s}^\circ. \tag{14}$$

Once the algorithm converges to or terminated by this stop criterion at the iteration k with $\mathbf{R}'$ near the global optimal $\mathbf{R}^*$, for any vector $\mathbf{v}$, we have

$$\begin{aligned} \theta(\mathbf{R}'\mathbf{v}, \mathbf{R}^*\mathbf{v}) &\leq \angle(\mathbf{R}'\mathbf{v}, \mathbf{R}^*\mathbf{v}) && \text{(by Lemma 3)} \\ &\leq d_\angle(\mathbf{R}', \mathbf{R}^*) && \text{(by Lemma 4)} \\ &\leq r_D \leq \sqrt{3}\sigma_k && \text{(by Corollary 1)} \\ &\leq 2/s^\circ. && \text{(by Eq. (14))} \end{aligned} \tag{15}$$

The proof will be the same flow for the azimuth criterion. We can choose ϵ to be $2/s^\circ$, where s is the parameter determined by user for the EGI resolution. Consequently, there exists a constant ϵ that guarantees ϵ-optimality. The proof is complete. □

REFERENCES

[1] J. M. Coughlan and A. L. Yuille, "Manhattan world: Compass direction from a single image by bayesian inference," in *IEEE International Conference on Computer Vision (ICCV)*, vol. 2, 1999, pp. 941–947. 1, 2

[2] J. Straub, G. Rosman, O. Freifeld, J. J. Leonard, and J. W. Fisher, "A mixture of manhattan frames: Beyond the manhattan world," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 3770–3777. 1, 2, 3, 7, 8

[3] B. Ghanem, A. Thabet, J. Carlos Niebles, and F. Caba Heilbron, "Robust manhattan frame estimation from a single rgb-d image," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 3772–3780. 1, 2, 7, 8

[4] J. Straub, N. Bhandari, J. J. Leonard, and J. W. Fisher III, "Real-time manhattan world rotation estimation in 3d," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 1913–1920. 1, 3, 7, 8

[5] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski, "Manhattan-world stereo," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 1422–1429. 1, 2

[6] S. N. Sinha, D. Steedly, and R. Szeliski, "Piecewise planar stereo for image-based rendering." in *IEEE International Conference on Computer Vision (ICCV)*, 2009, pp. 1881–1888. 1, 2

[7] V. Hedau, D. Hoiem, and D. Forsyth, "Recovering the spatial layout of cluttered rooms," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 1849–1856. 1, 2

[8] S. Ramalingam and M. Brand, "Lifting 3d manhattan lines from a single image," in *IEEE International Conference on Computer Vision (ICCV)*, 2013, pp. 497–504. 1

[9] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor segmentation and support inference from rgbd images," in *European Conference on Computer Vision (ECCV)*. Springer, 2012, pp. 746–760. 1, 2, 6, 8, 9, 10

[10] F. Endres, J. Hess, N. Engelhard, J. Sturm, D. Cremers, and W. Burgard, "An evaluation of the rgb-d slam system," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2012, pp. 1691–1696. 1

[11] W. Choi, Y.-W. Chao, C. Pantofaru, and S. Savarese, "Understanding indoor scenes using 3d geometric phrases," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 33–40. 1, 2

[12] S. Gupta, P. Arbelaez, and J. Malik, "Perceptual organization and recognition of indoor scenes from rgb-d images," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 564–571. 1, 2

[13] J.-C. Bazin, Y. Seo, R. Hartley, and M. Pollefeys, "Globally optimal inlier set maximization with unknown rotation and focal length," in *European Conference on Computer Vision (ECCV)*. Springer, 2014, pp. 803–817. 1, 2

[14] H. Zhou, D. Zou, L. Pei, R. Ying, P. Liu, and W. Yu, "Structslam: Visual slam with building structure lines," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 4, pp. 1364–1375, 2015. 1

[15] R. Horst and H. Tuy, *Global optimization: Deterministic approaches*. Springer Science & Business Media, 2013. 1, 3, 4, 6, 12

[16] A. J. Parra Bustos, T.-J. Chin, and D. Suter, "Fast rotation search with stereographic projections for 3d registration," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 3930–3937. 1, 2, 3

[17] B. K. Horn, "Extended gaussian images," *Proceedings of the IEEE*, vol. 72, no. 12, pp. 1671–1686, 1984. 1, 4

[18] K. Joo, T.-H. Oh, J. Kim, and I. S. Kweon, "Globally optimal manhattan frame estimation in real-time," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1

[19] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981. 1, 2

[20] J.-C. Bazin, Y. Seo, and M. Pollefeys, "Globally optimal consensus set maximization through rotation search," in *Asian Conference on Computer Vision (ACCV)*. Springer, 2012, pp. 539–551. 1, 2, 3, 5, 7, 8, 12

[21] D. C. Lee, M. Hebert, and T. Kanade, "Geometric reasoning for single image structure recovery," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2009, pp. 2136–2143. 2

[22] V. Hedau, D. Hoiem, and D. Forsyth, "Thinking inside the box: Using appearance models and context based on room geometry," in *European Conference on Computer Vision (ECCV)*. Springer, 2010, pp. 224–237. 2

[23] A. G. Schwing, T. Hazan, M. Pollefeys, and R. Urtasun, "Efficient structured prediction for 3d indoor scene understanding," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2012, pp. 2815–2822. 2

[24] L. D. Pero, J. Bowdish, D. Fried, B. Kermgard, E. Hartley, and K. Barnard, "Bayesian geometric modeling of indoor scenes," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 2719–2726. 2

[25] A. G. Schwing and R. Urtasun, "Efficient exact inference for 3d indoor scene understanding," in *European Conference on Computer Vision (ECCV)*. Springer, 2012, pp. 299–313. 2

[26] C. Taylor and A. Cowley, "Parsing indoor scenes using rgb-d imagery," in *Proceedings of Robotics: Science and Systems*, July 2012. 2, 8

[27] C. Zhang, L. Wang, and R. Yang, "Semantic segmentation of urban scenes using dense depth maps," in *European Conference on Computer Vision (ECCV)*. Springer, 2010, pp. 708–721. 2

[28] J.-C. Bazin, Y. Seo, C. Demonceaux, P. Vasseur, K. Ikeuchi, I. Kweon, and M. Pollefeys, "Globally optimal line clustering and vanishing point estimation in manhattan world," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 638–645. 2, 3, 10

[29] R. I. Hartley and F. Kahl, "Global optimization through rotation space search," *International Journal of Computer Vision (IJCV)*, vol. 82, no. 1, pp. 64–79, 2009. 2, 3, 5, 6, 11

[30] H. Li, "Consensus set maximization with guaranteed global optimality for robust geometry estimation," in *IEEE International Conference on Computer Vision (ICCV)*. IEEE, 2009, pp. 1074–1080. 3

[31] A. Makadia, A. Patterson, and K. Daniilidis, "Fully automatic registration of 3d point clouds," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, 2006, pp. 1297–1304. 4

[32] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, 2001, pp. I–511–I–518 vol.1. 5

[33] V. Balakrishnan, S. Boyd, and S. Balemi, "Branch and bound algorithm for computing the minimum stability degree of parameter-dependent linear systems," *International Journal of Robust and Nonlinear Control*, vol. 1, no. 4, pp. 295–317, 1991. 6, 12

[34] N. I. Fisher, *Statistical analysis of circular data*. Cambridge University Press, 1995. 6

[35] M. Grundmann, V. Kwatra, and I. Essa, "Auto-directed video stabilization with robust l1 optimal camera paths," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 225–232. 8, 10

[36] C. Jia and B. L. Evans, "Constrained 3d rotation smoothing via global manifold regression for video stabilization," *IEEE Transactions on Signal Processing*, vol. 62, no. 13, pp. 3293–3304, 2014. 8, 10

[37] P. Denis, J. H. Elder, and F. J. Estrada, *Efficient edge-based methods for estimating manhattan frames in urban imagery*. Springer, 2008. 10, 11

[38] E. Rosten and T. Drummond, "Fusing points and lines for high performance tracking," in *IEEE International Conference on Computer Vision (ICCV)*, vol. 2. IEEE, 2005, pp. 1508–1515. 8

[39] J. Shi and C. Tomasi, "Good features to track," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 1994, pp. 593–600. 8

[40] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, ISBN: 0521540518, 2004. 11